

Choice of boundary condition and collision operator for lattice-Boltzmann simulation of moderate Reynolds number flow in complex domains

Hywel B. Carver,^{1,2} Rupert W. Nash,¹ Miguel O. Bernabeu,^{1,2} James Hetherington,^{1,3} Derek Groen,¹ Timm Krüger,¹ and Peter V. Coveney^{1,*}

¹*Centre for Computational Science, Department of Chemistry, University College London, 20 Gordon Street, London, WC1H 0AJ, United Kingdom*

²*CoMPLEX, University College London, Physics Building, Gower Street, London, WC1E 6BT, United Kingdom*

³*Research Software Development Team, Research Computing and Facilitating Services, University College London, Podium Building - 1st Floor, Gower Street, London, WC1E 6BT, United Kingdom*

(Dated: November 2, 2012)

Modeling blood flow in larger vessels using lattice-Boltzmann methods comes with a challenging set of constraints: a complex geometry with walls and inlet/outlets at arbitrary orientations with respect to the lattice, intermediate Reynolds number, and unsteady flow. Simple bounce-back is one of the most commonly used, simplest, and most computationally efficient boundary conditions, but many others have been proposed. We implement three other methods applicable to complex geometries (Guo-Zheng-Shi, Bouzdi-Firdaouss-Lallemand, Junk-Yang) in our open-source application HEMELB. We use these to simulation Poiseuille and Womersley flow in a cylindrical pipe with an arbitrary orientation at physiologically relevant Reynolds (1–100) and Womersley (4–12) numbers and compare the accuracy to analytical solutions. We find that the Bouzdi-Firdaouss-Lallemand method offers the best accuracy and stability properties with first-order convergence in space. Simple bounce-back has similar behavior but with errors around 50% larger. The Guo-Zheng-Shi and Junk-Yang methods, while accurate at low Reynolds number, are unstable at Reynolds numbers ≥ 30 and so cannot be recommended for use in hemodynamic simulation of larger arteries. The choice of collision operator (lattice Bhatnagar-Gross-Krook vs. multiple relaxation time) and velocity set (D3Q15 vs. D3Q19) does not significantly affect the accuracy in the problems studied.

PACS numbers: 02.70.-c ; 47.11.-j ; 47.63.Cb

I. INTRODUCTION

In the last two decades, lattice-Boltzmann methods (LBM) [1–3] have been widely studied and used for fluid flow problems. We are particularly interested in its application to the study of hemodynamics: the flow behavior of blood under physiological conditions. Accurate simulation of the flow of blood in an individual could have near-term clinical benefits for, *inter alia*, the treatment of aneurysms [4–6] or stenoses [7–9].

Applications such as these have a number of challenges: retaining computational performance in a relatively sparse, three-dimensional (3D) fluid domain; capturing the complex rheology of a dense suspension of deformable particles; accounting for the compliance of vessel walls, and choosing a boundary condition method suitable for stable and accurate simulation. Indirect addressing [10, 11] offers a solution to the sparsity problem by precomputing the addresses of neighboring points instead of requiring that the points be stored in a dense, 3D array. We will not examine the rheology and compliance problems in this article, restricting ourselves to a Newtonian fluid within a rigid-walled system, but instead aim to provide recommendations for the choice of boundary condition method(s) for a complex geometry. This choice is examined in the context of two other variables:

the discrete velocity set and lattice-Boltzmann collision operator. Bearing in mind recent controversy over the reproducibility of computational science [12, 13], we have released the source code of our application HEMELB for the community’s scrutiny and use.

Lätt *et al.* [14] considered five boundary conditions and assess their accuracy. However, they studied only boundary conditions in which the wall passes through a lattice point, immediately restricting their results to boundaries that are normal to one of the Cartesian directions of the underlying grid. Boyd *et al.* [15] compared simple bounce-back to the Guo-Zheng-Shi method [16] in simulations of arterial bifurcations, finding that the two methods give different results when stenosis is present, but without any analysis suggesting which, if either, is more accurate. Stahl *et al.* [17] performed LBM simulations with simple bounce-back boundary conditions to examine the effect of staircased boundaries (where walls that are not aligned with the underlying grid are approximated by a set of grid edges) on the measurements of shear stress both at the wall and in the bulk flow, finding errors in the shear stress of up to 35% for two-dimensional (2D) channel flow and 28% for 2D bent-pipe flow. They also introduced a method for estimating the local wall normal from the shear stress tensor.

In this paper, we briefly introduce the lattice-Boltzmann method in Sec. II and discuss the collision operators, velocity sets and boundary conditions used. Our open-source software [18] and simulation protocols

* Corresponding author: p.v.coveney@ucl.ac.uk

are described, and results given in Sec. III. We present our conclusions in Sec. IV.

II. THE LATTICE-BOLTZMANN METHOD

Here we provide a brief summary of the lattice-Boltzmann method (LBM); for a full derivation, we refer the reader to one of the many thorough descriptions available [1–3]. LBM operates at a mesoscopic level, storing a discrete-velocity approximation to the one-particle distribution functions of the Boltzmann equation of kinetic theory [19], $\{f_i(\mathbf{r}, t)\}$, on a regular lattice, with grid spacing Δx . The set of velocities $\{\mathbf{c}_i\}$ is chosen such that the distances travelled in one timestep (Δt), $\mathbf{e}_i = \mathbf{c}_i \Delta t$, are lattice vectors and to ensure Galilean invariance [20]. When one only wishes to reproduce Navier-Stokes dynamics, the set is typically a subset of the Moore neighborhood, including the rest vector. For 3D simulations, the most commonly used sets have 15 and 19 members (termed D3Q15 and D3Q19, respectively). The LBM can be shown, through a Chapman-Enskog expansion, to reproduce the Navier-Stokes equations in the quasi-incompressible limit with errors proportional to the Mach number squared.

In the absence of forces, the density $\rho(\mathbf{r}, t)$ and the velocity $\mathbf{u}(\mathbf{r}, t)$ at a fluid site can be calculated from the distributions by

$$\rho = \sum_i f_i \quad (1)$$

$$\rho \mathbf{u} = \sum_i f_i \mathbf{c}_i. \quad (2)$$

Advancing the system one timestep is conceptually divided into two stages. The first is collision, which relaxes the distributions towards a local equilibrium (we denote the post-collisional distributions as f_i^*):

$$f_i^*(\mathbf{r}, t) = f_i(\mathbf{r}, t) + \hat{\Omega}(f_i(\mathbf{r}, t)), \quad (3)$$

where $\hat{\Omega}$ is the collision operator. The second is streaming, where the post-collisional distributions are propagated along the lattice vectors to new locations in the lattice, defining the distributions of the next timestep:

$$f_i(\mathbf{r} + \mathbf{c}_i \Delta t, t + \Delta t) = f_i^*(\mathbf{r}, t). \quad (4)$$

A. Collision operators

The collision operator approximates the microscopic interparticle interactions. Here we summarize the two considered in this article: lattice Bhatnagar-Gross-Krook and multiple relaxation time.

Lattice Bhatnagar-Gross-Krook (LBGK) [20, 21] approximates the collision process as relaxation towards

a local equilibrium (see below), in a discrete velocity analogue of the Boltzmann-BGK approximation from kinetic theory [22],

$$\hat{\Omega}(f_i) = -\frac{(f_i - f_i^{\text{eq}})}{\tau}, \quad (5)$$

where τ is the relaxation time. This can be shown, through a Chapman-Enskog expansion (see, e.g., [2, 23]), to reproduce the Navier-Stokes equations with a kinematic viscosity ν given by

$$\nu = c_s^2(\tau - \Delta t/2). \quad (6)$$

LBGK is simple to implement, gives excellent performance, and is therefore widely used.

Multiple relaxation time (MRT) collision operators, developed at the same time as LBGK [24], generalize the notion of relaxation towards local equilibrium by allowing different relaxation rates for different moments of the distributions, potentially improving stability properties [25]. The eigenvalue of the collision matrix which corresponds to the relaxation of shear stress λ_{shear} determines the viscosity as in (6) ($\tau \rightarrow 1/\lambda_{\text{shear}}$). We consider the MRT operator on the D3Q15 and D3Q19 lattices, as described by d’Humières *et al.* [25].

For the equilibrium distribution, we use a second-order (in velocity space) approximation to a Maxwellian distribution

$$f_i^{\text{eq}}(\rho, \mathbf{u}) = \rho w_i \left(1 + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{c}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2} \right), \quad (7)$$

where the weights w_i and speed of sound c_s depend on the choice of velocity set.

B. No-slip boundary conditions

Boundary conditions for LBM have some additional complications compared to boundary conditions for Navier-Stokes based methods, due to the LBM’s mesoscopic nature. Typically, one wishes to impose conditions on the macroscopic, hydrodynamic variables (ρ, \mathbf{u}) but these must be implemented through a closure relation for the mesoscopic distributions. There is no single, obviously superior choice. In this section, we will briefly review some commonly used boundary condition methods for lattice-Boltzmann models which impose the no-slip condition that the velocity of fluid adjacent to a wall is equal to the velocity of the wall.

Many boundary condition methods do not vary their behavior with respect to the location of the walls in relation to the Eulerian grid. The wall is often assumed to pass infinitesimally close to a point, the grid point remaining inside the fluid domain (sometimes referred to as “wet node” boundary conditions [14]), or alternatively the boundary is considered to be halfway along the lattice vector to a point outside the fluid. In the case of complex or non-lattice-aligned domains, these methods will

always cause a first-order *modeling* error, irrespective of the order of numerical accuracy of the resulting lattice-Boltzmann method [26]. This point has been studied numerically by Stahl *et al.* [17]. Other boundary conditions allow the wall to be at an arbitrary position along the link between a solid and a fluid site. These reduce the modeling error of fixed wall position methods, but often at the price of increased complexity and/or the requirement of data from neighboring fluid lattice points. Further, a number of methods suffer reduced accuracy at sites in corners, which can reduce the accuracy of simulations throughout the domain.

Simple bounce-back (SBB) is perhaps the most widely used boundary condition for solid walls, positioning them halfway along the lattice vector from fluid to solid. It is straightforward to implement and gives second-order accurate simulation results for flat boundaries aligned with the Cartesian axes of the lattice [27], although in more complex cases this degrades to first-order accuracy [28]. It is also computationally cheap and can be local in its operation. SBB ensures conservation of mass up to machine precision. In this work we use the halfway bounce-back scheme.

There are a number of popular methods [29–32] that can only operate on straight, axis-aligned planes and force the boundary to pass directly through the lattice point. Malaspinas *et al.* [33] generalized the regularized method [31] to cope accurately with corner nodes, although the authors acknowledge that it fails for the case of a D3Q15 lattice and a right-angled corner. These methods are, therefore, unsuitable for problems involving complex boundaries.

The **Bouzidi-Firadaouss-Lallemand (BFL)** method [34] starts with simple bounce-back, but interpolates the value of the to-be-propagated distribution with the distribution at the fluid site which standard bounce-back would stream it to. They present two variants: one using linear interpolation and another using quadratic interpolation. In the present work, we restrict our attention to the linear case, due to its locality and smaller communication requirement (indeed, it can be implemented without any inter-process communication above the normal lattice-Boltzmann streaming step, albeit at a price of revisiting the sites adjacent to the wall); the authors claim that both variants show second-order convergence, but with the linear method having a poorer prefactor [34]. However inspection of Fig. 6 in [34] appears to show first-order convergence, at least for the largest resolution results shown.

Guo, Zheng, and Shi (GZS) [16] present a boundary condition which decomposes the unknown distributions at the wall into equilibrium and non-equilibrium parts. The equilibrium part uses the density of the fluid site and a linearly extrapolated velocity such that the velocity at the solid wall is as imposed. For the non-equilibrium part, the value from the fluid site (or the next site into the bulk) is used.

Junk and Yang (JY) [35] propose a correction to

the simple bounce-back method. They claim an advantage compared to interpolation-based methods such as GZS and BFL as their method is completely local and is able to handle non-straight boundaries where sites have lattice vectors in opposite directions, both crossing the solid-fluid boundary. The method arises from their analysis [36] of boundary conditions for LBM in terms of general methods for studying finite difference schemes rather than the standard Chapman-Enskog expansion. By adding correction terms to the collision operator and then discretizing them in an optimal (under their analysis framework) manner, they ensure the Navier-Stokes equations are obeyed with the correct boundary conditions.

We also note that some authors propose the use of grid-refinement [37], finite difference [38], and finite volume [39, 40] discretizations of the discrete Boltzmann equation as methods for improving accuracy around non-planar boundaries, but we will restrict our attention here to implementations on a single, regular grid. Additionally, the immersed boundary method (IBM) [41] has been used in conjunction with the LBM to simulate rigid [42] and deformable [43] boundaries. IBM requires a further layer of fluid sites outside the walls as well as another, simpler boundary condition at the edge of the halo region, but admits an extension to moving boundaries in a natural way.

C. Open boundary conditions

For hemodynamic applications, we must be able to specify inlet and outlet boundary conditions. If the velocity and pressure at the open boundary is known, then one can choose Dirichlet boundary conditions. However, if one wishes to use Dirichlet conditions at both inlets and outlets, the total mass in the system can increase or decrease unboundedly unless the mass fluxes into and out from the system balance exactly [44]. Alternatively, one could impose mixed Dirichlet-Neumann boundary conditions:

$$p = p_0, \quad (8a)$$

$$\mathbf{u}_{\parallel} = 0, \quad (8b)$$

$$\hat{\mathbf{n}} \cdot \nabla u_{\perp} = 0, \quad (8c)$$

where $\hat{\mathbf{n}}$ is the inward pointing normal of the open boundary.

A number of authors [30, 45–47] have proposed open boundary conditions that fulfill some or all of these requirements, however the techniques cited are only suitable for inlets aligned with the lattice's axes. We have therefore developed a simple method that meets these requirements.

Assume that, at the start of a timestep, all distributions for an inlet/outlet site are known. LBM then proceeds as normal: a collision step, followed by a streaming

step; the distributions that would have come from exterior sites now have an undefined value. We close the system by constructing a “phantom site” (indicated below with a subscript \mathcal{P}) beyond the boundary, whose hydrodynamic variables are estimated based on the imposed values and those at the inlet/outlet site (indicated with subscript \mathcal{I}). Note that there is one phantom site per missing distribution, i.e., the phantom sites of adjacent inlet/outlet sites are unrelated. This is to eliminate the need for communication between neighboring sites. The equilibrium distribution for the missing direction is computed at the phantom site and then streamed into place. For the density, we assume the target pressure p_0 , i.e., we perform a zero-order extrapolation from the inlet plane. Condition (8b) is enforced by projecting away any velocity component not parallel to the inlet/outlet normal, $\hat{\mathbf{n}}$. For condition (8c), we take first-order finite-difference approximations for the derivatives, giving:

$$\hat{\mathbf{n}} \cdot \nabla u_{\perp} \approx \frac{u_{\perp}(\mathbf{r}_{\mathcal{I}}, t) - u_{\perp}(\mathbf{r}_{\mathcal{P}}, t)}{c_{i\alpha} \cdot \hat{\mathbf{n}} \Delta t} = 0. \quad (9)$$

Hence

$$\mathbf{u}(\mathbf{r}_{\mathcal{P}}, t) \approx (\mathbf{u}(\mathbf{r}_{\mathcal{I}}, t) \cdot \hat{\mathbf{n}}) \hat{\mathbf{n}}. \quad (10)$$

For lattice sites that are adjacent to both open and closed boundaries, SBB is applied to the distributions in the direction of the wall. In the data presented in Sec. III B, we see that this boundary condition gives good agreement with the expected velocity fields, but in simulations at higher Reynolds number (data not shown) we observe larger errors, especially at these boundaries. This will be explored further in subsequent work.

III. SIMULATIONS

Our goal is to determine which combination of boundary condition, collision operator and velocity set gives the best all-round accuracy in a general, complex geometry, with a focus on computational hemodynamics. We also assess the computational requirements of the different models.

We compare against analytical solutions, which restricts us to relatively simple domains: we choose a cylinder with both steady, pressure-driven flow (Hagen-Poiseuille flow) and with a time-dependent, sinusoidal pressure gradient (Womersley flow). By choosing a non-axis-aligned orientation for the cylinder, we better mimic a typical production simulation of the human vasculature. The orientation $\hat{\mathbf{n}}$ was chosen pseudorandomly from the unit sphere, subject to the constraint that $\hat{\mathbf{n}} \cdot \hat{\mathbf{e}}_i \leq 0.9$, $\forall i$. The value is

$$\hat{\mathbf{n}} = [-0.299, 0.382, 0.874]. \quad (11)$$

Our approach is to select parameters in lattice units, but with physiologically relevant Reynolds (Re) and Womersley (α) numbers.

A. Software: HemeLB

The simulations in this paper were performed with HEMELB [11], a lattice-Boltzmann-based fluids solver, which includes capability for *in situ* imaging of flow-fields and real-time steering [48]. It is a distributed memory application, parallelized with MPI. We have shown that HEMELB’s computational performance scales linearly up to at least 32,768 cores [49]. We have released the software online [18], under the open-source GNU Lesser General Public License (LGPL), to enable interested researchers to reproduce our results as well as to use the software for novel problems.

HEMELB has several linked components, described in Groen *et al.* [49]. We have recently re-developed the lattice-Boltzmann core to allow for easy switching between use of different velocity sets, collision operators, and boundary conditions, through a statically polymorphic, object-oriented design. This avoids any run-time overhead due to dynamic polymorphism. The individual software components are tested through a battery of over one hundred unit and regression tests, which are run nightly by our continuous integration server.

HEMELB includes, amongst other features: the D3Q15 and D3Q19 velocity sets; the lattice Bhatnagar-Gross-Krook (LBGK) and multiple relaxation time (MRT) collision operators; and, the simple bounce-back (SBB), Guo-Zheng-Shi (GZS), Bouzidi-Firdaouss-Lallemand (BFL), and Junk-Yang (JY) boundary condition methods.

The software includes a separate tool for defining the simulation domain. This requires either a geometric primitive or a general surface, meshed with triangles. The user can then place inlets and outlets, specify their pressure conditions, and select the fineness of the lattice. This setup tool then generates the input for HEMELB itself, producing a description of each fluid site and, if needed, the location of the wall.

B. Convergence analysis

In this section we report on a series of simulations of Hagen-Poiseuille flow over a range of resolutions and Reynolds numbers, defined here as

$$\text{Re} = U_{\max} D / \nu, \quad (12)$$

where U_{\max} is the maximum velocity, D the pipe diameter, and ν the kinematic viscosity. The velocity \mathbf{U} is

$$\mathbf{U} = \frac{\nabla p}{4\rho\nu} ((D/2)^2 - r^2) \hat{\mathbf{n}}, \quad (13)$$

where r is the distance from the cylinder axis, defined by $\hat{\mathbf{n}}$, and ∇p is the pressure gradient.

In Table I we list all the parameters chosen for the simulations. The range of Re spans typical values for cerebral arteries in the human body [50]. For each case

TABLE I. Parameters for convergence analysis. All values are given in lattice units. Parameters are, from left to right: the Reynolds number Re ; the cylinder diameter D ; the predicted Mach number Ma ; the LBGK relaxation time τ ; the relative density difference imposed $\Delta\rho/\rho_0$; the momentum diffusion time $T_{\text{mom}} \equiv D^2/\nu$, and the sound propagation time $T_s \equiv L/c_s$.

Re	D	Ma	τ	$\Delta\rho/\rho_0$	T_{mom}	T_s
1	12	0.0241	1.00	0.01850	864	42
1	24	0.0120	1.00	0.00463	3460	83
1	48	0.0060	1.00	0.00116	13800	166
1	96	0.0030	1.00	0.00029	55300	333
1	192	0.0015	1.00	0.00007	221000	665
30	12	0.3610	0.75	0.13900	1730	42
30	24	0.1800	0.75	0.03470	6910	83
30	48	0.0902	0.75	0.00868	27600	166
30	96	0.0451	0.75	0.00217	111000	333
30	192	0.0226	0.75	0.00054	442000	665
100	12	0.4810	0.60	0.07410	4320	42
100	24	0.2410	0.60	0.01850	17300	83
100	48	0.1200	0.60	0.00463	69100	166
100	96	0.0601	0.60	0.00116	276000	333
100	192	0.0301	0.60	0.00029	1110000	665

we vary the diameter D from 12 – 192 lattice units; the length of tube used is given by $L = 2D$. Due to the finite speed of sound in LBM ($c_s = 1/\sqrt{3}$ for the models used here), we list the Mach numbers ($Ma \equiv U_{\text{max}}/c_s$); the lowest resolution simulations in each have extremely high values and will consequently have poor accuracy, but this allows us to assess the convergence behavior at modest computational expense. Next we show the value of the LBGK relaxation time τ which must be greater than $1/2$ [1] and not be much greater than one [51]. We hold τ constant in lattice units while refining the spatial resolution, which implies diffusive scaling of the timestep (when converted to physical units), i.e., $\Delta t \propto \Delta x^2$. The density, and hence pressure, difference $\Delta\rho$ driving the flow is also reported; this must remain much less than the reference density of the simulation ρ_0 . Finally, we list the time for momentum to diffuse across the cylinder's diameter, $T_{\text{mom}} \equiv D^2/\nu$, and the time for a sound wave to propagate the length of the cylinder, $T_s \equiv L/c_s$, to give some idea of the time required for the simulation to converge to a steady state. To determine whether a simulation has indeed converged, we compute the maximum difference between flow fields at two times

$$\Delta u(t_1, t_2) = \frac{\max_{\mathbf{r}} \|\mathbf{u}(\mathbf{r}, t_1) - \mathbf{u}(\mathbf{r}, t_2)\|}{U_{\text{max}}}, \quad (14)$$

and require that $\Delta u(t, t+1) < 10^{-7}$.

We use a simple initialization procedure, initializing to a uniform density fluid at rest. This approach is general and can be applied to any geometry without requiring any preprocessing step [52–54], but does require longer simulation times until a steady solution is reached. For each simulation we compare the Poiseuille solution, $\mathbf{U}(\mathbf{r})$,

with the velocity field found by simulation, $\mathbf{u}(\mathbf{r}, t)$. We define the velocity error as

$$\mathbf{u}^*(\mathbf{r}, t) \equiv \mathbf{u}(\mathbf{r}, t) - \mathbf{U}(\mathbf{r}) \quad (15)$$

and use the ℓ^2 -norm scaled by the predicted velocity range as our measure of error:

$$\epsilon_u^2(t) = \frac{\sqrt{\sum_{\mathbf{r}} \mathbf{u}^{*2}}}{\sqrt{N} \max_{\mathbf{r}} \|\mathbf{U}\|}. \quad (16)$$

These are evaluated over all fluid sites in the central 90% of the cylinder, thus excluding the inlet and outlet sites. For the pressure gradient, we use the measured difference at the edge of this volume divided by the distance between the two planes. This is to disentangle the errors due to the open boundary condition method from the no-slip condition. We hope to return in the future to a full validation of the open boundary condition method.

The simulations were performed on HECToR, the UK national supercomputer, using between 32 and 16384 cores. The number of cores for each run was chosen to minimize the run time while remaining efficient which we have shown to occur at around 10^3 sites per core [49]; we take the integer power of two that is closest to this.

In Fig. 1 we show the velocity errors as a function of increasing resolution for each of the four boundary condition methods and the two velocity sets. We use only LBGK as the collision operator. We note that many of the simulations using the Guo-Zheng-Shi and Junk-Yang boundary conditions became unstable during the simulation, with the lattice-Boltzmann distributions becoming negative. These methods only ran until a converged solution was obtained for the $Re = 1$ case.

The Junk-Yang method with a D3Q15 velocity set showed particularly poor stability, only converging for the lowest resolution case and to a poor approximation of the expected flow. Our unit test suite demonstrates (data not shown) that the JY implementation reduces to SBB when the wall is a plane normal to one of the axes and halfway between the site and its non-fluid neighbors, as expected [35]. In the D3Q19 case, JY gives better accuracy than SBB but poorer than BFL. The GZS method on D3Q19 also had poor stability, failing for nearly all the simulations, while GZS with D3Q15 gave good results for the lower resolution $Re = 1$ simulations.

Simple bounce-back and Bouzidi-Firdaouss-Lallemand were both reliable, remaining stable at all the reported Reynolds number and resolutions, except for BFL with D3Q15 at the highest resolutions. These both show first-order convergence to the analytical solution as the space step is reduced, as expected. BFL shows better agreement with the solutions than SBB, having errors around one third lower across all simulations. As the Reynolds number is increased, we note that the errors also increase, particularly when going from $Re = 30$ to 100. This may be due to the combination of the increase in Mach number and the decrease in τ towards the lower limit of 0.5 [51].

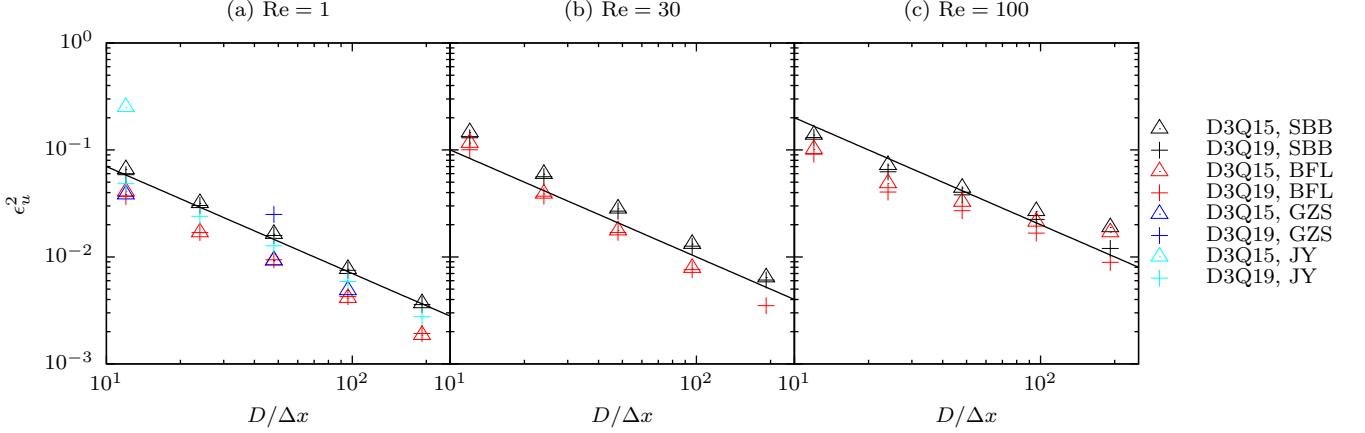


FIG. 1. (Color online.) Convergence of velocity error residuals (ϵ_u^2) as defined in Eq. (16) vs. diameter in lattice units. (a) Reynolds number $Re = 1$; (b) $Re = 30$; (c) $Re = 100$. Colors (black, red, blue, cyan) indicate boundary condition (SBB, BFL, GZS, JY), triangles indicate D3Q15 and crosses D3Q19. The solid lines are guides to the eye showing first-order convergence. Missing symbols are due to simulations becoming unstable for that combination of LBM and parameters.

C. Womersley flow

Here we report on simulations of oscillatory flow, in order to explore the different boundary condition methods' effect on accuracy for time-dependent cases. The Womersley number (α) is a dimensionless number governing dynamical similarity in cases of oscillatory flow. It relates to the ratio of transient forces to viscous forces (or alternatively the ratio of diameter to boundary layer growth during one period of oscillation) and is defined as:

$$\alpha = L \sqrt{\frac{\omega}{\nu}}, \quad (17)$$

where ω is the characteristic angular frequency of the pressure oscillation and L is a characteristic length scale.

In the case of a cylinder (radius R , axis $\hat{\mathbf{n}}$) and laminar flow with zero average pressure gradient ($\nabla p(t) = (A/L) \sin \omega t$, $\omega \equiv 2\pi/T_{\text{osc}}$), the time-dependent Navier-Stokes equations admit an analytic solution [55]:

$$\mathbf{u}(r, t) = -\Re \left[\frac{A}{\omega \rho L} \left(1 - \frac{J_0(i^{3/2} \alpha \frac{r}{R})}{J_0(i^{3/2} \alpha)} \right) e^{i\omega t} \right] \hat{\mathbf{n}}, \quad (18)$$

where J_0 is the order-0 Bessel function of the first kind and α is defined using the following characteristic quantities: cylinder radius, the pressure oscillation angular frequency, and the fluid viscosity.

In the larger arteries of the human body, Womersley numbers range approximately between 4 and 20 [50], so we select from this range. For these simulations, we define the Reynolds number as that for a Poiseuille flow with a pressure gradient given by the amplitude of the imposed gradient. We note that in the body, the Reynolds number scales approximately with α , so we select parameters only from a single curve within the $Re - \alpha$ plane, as

TABLE II. Parameters for Womersley flow simulations. All values are given in lattice units. Parameters are: Reynolds number (defined for a Poiseuille flow with the maximum pressure gradient); Womersley number; predicted Mach number (based on the maximum Poiseuille flow); the LBGK relaxation time; the maximum relative density difference; the period of the pressure oscillation, and the momentum diffusion time.

Re	α	Ma	τ	$\Delta\rho/\rho_0$	T_{osc}	T_{mom}	T_s
30	4	0.043	0.620	0.00200	5655	57600	166
100	8	0.078	0.565	0.00194	2618	107000	166
300	12	0.113	0.531	0.00135	2417	222000	166

shown in Table II. We use the same cylinder orientation as in Sec. III B and select a diameter $D = 48\Delta x$, as this gives a reasonable balance between computational cost and accuracy, such as would be chosen for production simulations; we keep $L = 2D$. We initialize the simulation to a constant pressure with zero velocity (with the phase of the pressure oscillation chosen such that the driving difference is zero at simulation start). We allow the simulation to run until $\Delta u(t, t - T_{\text{osc}}) < 10^{-7}$ is reached for all sample points during one oscillation; however, to reduce the amount of data collected, we record data only for those points within one lattice unit of an axis-normal plane halfway along the cylinder.

We run these simulations for *all* combinations of LBM components and compute residuals at four sample points during one pressure oscillation period using Eq. (16). The four residuals are then reduced by taking the root-mean-square average and the maximum respectively, effectively extending the averaging/maximization over time as well as space.

We found that all the simulations with the Guo-Zheng-Shi and Junk-Yang boundary conditions became unsta-

TABLE III. Error residuals for pulsatile flow simulations with different LBM components. On the left are the boundary condition method (BC); collision operator (CO); and velocity set (DmQn). On the right are the error residuals ϵ_u^2 for different values of Womersley number α (other parameters are as shown in Table II).

BC	CO	DmQn	ϵ_u^2		
			$\alpha = 4$	$\alpha = 8$	$\alpha = 12$
SBB	LBGK	D3Q15	0.025	0.043	0.061
BFL	LBGK	D3Q15	0.018	0.027	0.041
SBB	LBGK	D3Q19	0.026	0.043	0.060
BFL	LBGK	D3Q19	0.020	0.028	0.040
SBB	MRT	D3Q15	0.024	0.043	0.065
BFL	MRT	D3Q15	0.018	0.027	0.041
SBB	MRT	D3Q19	0.027	0.046	0.069
BFL	MRT	D3Q19	0.020	0.028	0.041

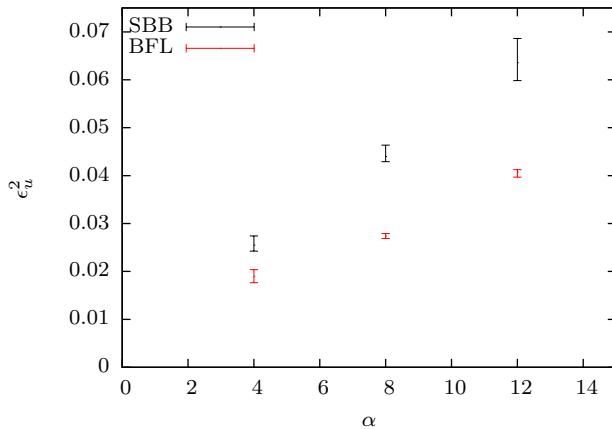


FIG. 2. (Color online.) Error residuals ϵ_u^2 as defined in Eq. (16) for simulations of pulsatile flow vs. Womersley number α for the simple bounce-back (SBB, black) and Bouzidi-Firdaouss-Lallemand (BFL, red) boundary conditions. Bars show the range of error residuals measured across all combinations of collision operator and velocity set.

ble before the end of the first period, as might have been expected since they failed for all the Poiseuille simulations with $\text{Re} \geq 30$ in Sec. III B. The results of the simple bounce-back and Bouzidi-Firdaouss-Lallemand simulations are shown in Table III. We see that all simulations well approximate the analytical solutions, with errors in the range 1.8–6.9%.

In Fig. 2, we show the range of error residuals for SBB and BFL against Womersley number. This clearly shows the BFL method offers superior accuracy to SBB in all cases, irrespective of the choice of collision operator and velocity set. Choosing MRT over LBGK does not offer any benefits, however we have not varied the relaxation rates for the non-stress tensor moments of the distributions (e.g., by projecting out all the “ghost” kinetic modes every timestep). We note again that we have adopted the parameters from [25] without optimization.

TABLE IV. Performance per core for different combinations of boundary condition, collision operator and velocity set. The columns indicate the combination of collision operator and velocity set, while the rows indicate the boundary condition. Unstable, and hence aborted, simulations are indicated with a dash (-). The values are given in millions of site updates per second (MSUPS) followed by the time to perform one boundary condition site update, relative to SBB, in parentheses.

	LBGK				MRT			
	D3Q15		D3Q19		D3Q15		D3Q19	
	SBB	1.86	(1)	1.59	(1)	0.67	(1)	0.45
BFL	1.70	(2.0)	1.46	(2.0)	0.61	(2.1)	0.42	(1.8)
GZS	1.40	(4.7)	-	-	0.60	(2.3)	0.40	(2.4)
JY	0.54	(28)	-	-	0.35	(11)	-	-

The choice of velocity set does not make a significant difference to the accuracy.

D. Relative performance

The relative performance of the options is germane to the choice of which combination of lattice-Boltzmann components to use. In Table IV we give the measured per-core performance for the lowest Re/α pulsatile flow simulations. All these runs were performed on HEC-ToR, a Cray XE6 supercomputer with two 16-core AMD Opteron 2.3 GHz Interlagos processors per node. The simulations used 256 cores. The performance is given in millions of site updates per second (MSUPS) and is based on timings that include only the lattice-Boltzmann updates. We also estimate the relative time to perform one site update for the different methods, assuming that an SBB site update takes the same time as a bulk fluid site update. The simulation domain includes 173,706 sites of which 9% are at the solid-fluid boundary and hence use the various boundary condition methods.

We do not place much emphasis on these results as we have not significantly optimized the implementation of boundary conditions (other than SBB) and the MRT collision operator is implemented naïvely. Nonetheless, simple bounce-back gives the best computational performance in all cases while BFL only needs approximately twice the compute time across the different simulations. The Junk-Yang implementation requires a number of linear algebra operations at every timestep (implemented in HEMELB with the Boost uBLAS library), explaining its relatively poor performance.

IV. CONCLUSION

The majority of benchmark problems reported in the lattice-Boltzmann literature use lattice-aligned geometries, rather than the complex domains required by many applications. We have performed a comparison between

LBM simulation solutions, from our open source application HEMELB [18], and analytical solutions in a non-lattice-aligned, curved domain up to Reynolds numbers of 100, with steady and unsteady flow. We have varied the resolution of the grid used and the different components of the algorithm (collision operator, velocity set, no-slip boundary condition). We find that at these moderate values of Re , the choice of velocity set and collision operator do not greatly affect accuracy or stability, but that the choice of no-slip boundary condition method is critical.

The Guo-Zheng-Shi and Junk-Yang methods, while offering good accuracy for steady flow at low Reynolds number, show poor stability at larger Re . They are therefore unsuitable for our hemodynamic applications or other applications that require even moderate Reynolds numbers.

We find that simple bounce-back (SBB) and the Bouzidi-Firdaouss-Lallemand (BFL) methods both show first-order convergence over a wide range of resolutions and Reynolds numbers, but with SBB having errors about 50% larger. The BFL algorithm is a factor of two slower than SBB but, since it is only used at the boundaries, it will cause an increase in CPU time by a factor of $\lesssim 1.1$. Reaching equivalent accuracy with SBB requires

enhancing the resolution by a factor of ~ 1.5 ; due to the diffusive scaling of the timestep, this increases CPU time by a factor of $\sim 1.5^5 \approx 7.5$, as well as increasing total memory consumption by a factor of $\sim 1.5^3 \approx 3$. We therefore recommend BFL in preference to SBB. Although SBB is less accurate, it is still acceptable for use in less demanding applications and when development time is at a premium. We hope that these results will prove helpful to the community when selecting methods for simulating hemodynamics and comparable applications with lattice-Boltzmann methods.

ACKNOWLEDGMENTS

This work made use of HECToR, the UKs national high-performance computing service. We also thank Dr Jörg Saßmannshausen for local computational support. This work was supported by the British Heart Foundation, EPSRC grants “Large Scale Lattice Boltzmann for Biocolloidal Systems” (EP/I034602/1) and 2020 Science (<http://www.2020science.net/>, EP/I017909/1), and the EC-FP7 projects CRESTA (<http://www.cresta-project.eu/>, grant no. 287703) and MAPPER (<http://www.mapper-project.eu/>, grant no. 261507).

[1] S. Succi, *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond* (Clarendon Press, Oxford, 2001).

[2] S. Chen and G. Doolen, *Annu. Rev. Fluid Mech.*, **30**, 329 (1998).

[3] C. K. Aidun and J. R. Clausen, *Annu. Rev. Fluid Mech.*, **42**, 439 (2010).

[4] E. O. Kung, A. S. Les, F. Medina, R. B. Wicker, M. V. McConnell, and C. A. Taylor, *J. Biomech. Eng.*, **133**, 041003 (2011).

[5] M. C. Villa-Uriol, I. Larrabide, J. M. Pozo, M. Kim, O. Camara, M. De Craene, C. Zhang, A. J. Geers, H. Morales, H. Bogunović, R. Cardenes, and A. F. Frangi, *Phil. Trans. R. Soc. A*, **368**, 2961 (2010).

[6] X. He, G. Duckwiler, and D. Valentino, *Comput. Fluids*, **38**, 789 (2008).

[7] R. L. Spilker, J. A. Feinstein, D. W. Parker, V. M. Reddy, and C. A. Taylor, *Ann. Biomed. Eng.*, **35**, 546 (2007).

[8] C. A. C. Taylor and C. A. C. Figueroa, *Annu Rev Biomed Eng*, **11**, 109 (2009).

[9] H. Tahir, A. G. Hoekstra, E. Lorenz, P. V. Lawford, D. R. Hose, J. Gunn, and D. J. W. Evans, *Interface Focus*, **1**, 365 (2011).

[10] M. Schultz, M. Krafczyk, J. Tölke, and E. Rank, in *High-Performance Scientific and Engineering Computing, Proceedings of the 3rd International FORTWIHR Conference on HPSEC, Erlangen, March 12–14, 2001*, edited by M. Breuer, F. Durst, and C. Zenger (Springer, Berlin, 2002) pp. 115–122.

[11] M. D. Mazzeo and P. V. Coveney, *Comput. Phys. Commun.*, **178**, 894 (2008).

[12] Z. Merali, *Nature*, **467**, 775 (2010).

[13] A. Morin, J. Urban, P. D. Adams, I. Foster, A. Sali, D. Baker, and P. Sliz, *Science*, **336**, 159 (2012).

[14] J. Lätt, B. Chopard, O. Malaspinas, M. Deville, and A. Michler, *Phys. Rev. E*, **77**, 056703 (2008).

[15] J. Boyd, J. M. Buick, J. Cosgrove, and P. Stansell, *Australasian Physical and Engineering Sciences in Medicine*, **27**, 207 (2004).

[16] Z. Guo, C. Zheng, and B. Shi, *Phys. Fluids*, **14**, 2007 (2002).

[17] B. Stahl, B. Chopard, and J. Lätt, *Comput. Fluids*, **39**, 1625 (2010).

[18] <http://ccs.chem.ucl.ac.uk/hemelb>.

[19] R. L. Liboff, *Kinetic Theory: Classical, Quantum, and Relativistic Descriptions*, 2nd ed. (John Wiley & Sons, Ltd., New York, 1998) ISBN 0471152986.

[20] Y. H. Qian, D. d’Humières, and P. Lallemand, *Europhys. Lett.*, **17**, 479 (1992).

[21] H. Chen, S. Chen, and W. H. Matthaeus, *Phys. Rev. A*, **45**, R5339 (1992).

[22] Bhatnagar, E. P. Gross, and M. Krook, *Phys. Rev.*, **94**, 511 (1954).

[23] A. J. C. Ladd and R. Verberg, *J. Stat. Phys.*, **104**, 1191 (2001).

[24] D. d’Humières, *Progress in Aeronautics and Astronautics*, **159**, 450 (1992).

[25] D. d’Humières, I. Ginzburg, M. Krafczyk, P. Lallemand, and L.-S. Luo, *Phil. Trans. R. Soc. A*, **360**, 437 (2002).

[26] Typically second order accuracy is sought, since this is the inherent accuracy of standard lattice-Boltzmann methods in bulk fluids.

[27] D. P. Ziegler, *J. Stat. Phys.*, **71**, 1171 (1993).

[28] I. Ginzburg and P. M. Adler, *J. Phys. II France*, **4**, 191 (1994).

- [29] T. Inamuro, M. Yoshino, and F. Ogino, *Phys. Fluids*, **7**, 2928 (1995).
- [30] Q. Zou and X. He, *Phys. Fluids*, **9**, 1591 (1997).
- [31] J. Lätt, *Hydrodynamic limit of lattice Boltzmann equations*, Ph.D. thesis, Université de Genève (2007).
- [32] S. Asumali and I. V. Karlin, *Phys. Rev. E*, **66**, 026311 (2002).
- [33] O. Malaspinas, B. Chopard, and J. Latt, *Comput. Fluids*, **49**, 29 (2011).
- [34] M. Bouzidi, M. Firdaouss, and P. Lallemand, *Phys. Fluids*, **13**, 3452 (2001).
- [35] M. Junk and Z. Yang, *Phys. Rev. E*, **72**, 066701 (2005).
- [36] M. Junk and Z. Yang, *J. Stat. Phys.*, **121**, 3 (2005).
- [37] O. Filippova and D. Hänel, *J. Comp. Phys.*, **147**, 219 (1998).
- [38] X. He, L.-S. Luo, and M. Dembo, *J. Comp. Phys.*, **129**, 357 (1996).
- [39] H. Chen, *Phys. Rev. E*, **58**, 3955 (1998).
- [40] S. Ubertini, G. Bella, and S. Succi, *Phys. Rev. E*, **68**, 016701 (2003).
- [41] C. S. Peskin, *Acta Numerica*, **11**, 479 (2002).
- [42] Z.-G. Feng and E. Michaelides, *J. Comp. Phys.*, **195**, 602 (2004).
- [43] J. Zhang, P. C. Johnson, and A. S. Popel, *Phys. Biol.*, **4**, 285 (2007).
- [44] T. Krüger, F. Varnik, and D. Raabe, *Phys. Rev. E*, **79**, 046704 (2009).
- [45] S. S. Chikatamarla, S. Asumali, and I. V. Karlin, *Eur. Phys. Lett.*, **74**, 215 (2006).
- [46] K. Mattila, J. Hyvälä, A. A. Folarin, and T. Rossi, *Int. J. Numer. Meth. Fluids*, **63**, 638 (2010).
- [47] K. Mattila, J. Hyvälä, and T. Rossi, *J. Stat. Mech.*, **2009**, P06015 (2009).
- [48] M. D. Mazzeo, S. Manos, and P. V. Coveney, *Comput. Phys. Commun.*, **181**, 355 (2010).
- [49] D. Groen, J. Hetherington, H. B. Carver, R. W. Nash, M. O. Bernabeu, and P. V. Coveney, “Analyzing and modeling the performance of the HemeLB lattice-Boltzmann simulation environment,” (2012), arXiv:1209.3972v1 [cs.PF].
- [50] J. D. Bronzino, *Biomedical Engineering Fundamentals (The Biomedical Engineering Handbook)*, 3rd ed. (CRC Press, 2006) ISBN 0849321212.
- [51] D. J. Holdych, D. R. Noble, J. G. Georgiadis, and R. O. Buckius, *J. Comp. Phys.*, **193**, 595 (2004).
- [52] A. Caiazzo, *J. Stat. Phys.*, **121**, 37 (2005).
- [53] A. M. Artoli, A. G. Hoekstra, and P. M. A. Sloot, *Comput. Fluids*, **35**, 227 (2006).
- [54] R. Mei, L. S. Luo, P. Lallemand, and D. d’Humières, *Comput. Fluids*, **35**, 855 (2006).
- [55] L. Formaggia, A. Quarteroni, and A. Veneziani, *Cardiovascular Mathematics* (Springer-Verlag Italia, Milano, 2009) ISBN 9788847011519.